

**MAINTAINING APPLICATION OPERATIONS WITHIN A SUBOPTIMAL GRID
ENVIRONMENT**

CROSS-REFERENCE TO RELATED APPLICATIONS

5

The present application is related to the following co-pending applications:

(1) U.S. Patent Application Serial No. ____/____ (Attorney Docket No.

AUS920030857US1).

10

BACKGROUND OF THE INVENTION

1. Technical Field:

15 **[0001]** The present invention relates in general to grid environments and in particular to managing application operation in a grid environment. Still more particularly, the present invention relates to maintaining application operation within a suboptimal grid environment by reconfiguring the application according to an application profile which expresses the operational requirements of an application in a grid environment.

20

2. Description of the Related Art:

[0002] Ever since the first connection was made between two computer systems, new ways of transferring data, resources, and other information between two computer systems via a connection continue to develop. In a typical network architecture, when two computer systems are exchanging data via a connection, one of the computer systems is considered a client sending requests and the other is considered a server processing the requests and returning results. In an effort to increase the speed at which requests are handled, server systems continue to expand in size and speed. Further, in an effort to handle peak periods when multiple requests are arriving every second, server systems are often joined together as a group and requests are distributed among the grouped servers. Multiple methods of grouping servers have developed such as clustering, multi-system shared data (sysplex) environments, and enterprise systems. With a cluster of servers, one server is typically designated to manage distribution of incoming requests and outgoing responses. The other servers typically operate in parallel to handle the distributed requests from clients. Thus, one of multiple servers in a cluster may service a client request without the client detecting that a cluster of servers is processing the request.

[0003] Typically, servers or groups of servers operate on a particular network platform, such as Unix or some variation of Unix, and provide a hosting environment for running applications. Each network platform may provide functions ranging from database integration, clustering services, and security to workload management and problem determination. Each network platform typically offers different implementations, semantic behaviors, and application programming interfaces (APIs).

[0004] Merely grouping servers together to expand processing power, however, is a limited method of improving efficiency of response times in a network. Thus, increasingly, within a company network, rather than just grouping servers, servers and groups of server systems are organized as distributed resources. There is an increased effort to collaborate, share data, share cycles, and improve other modes of interaction among servers within a company network and outside the company network. Further, there is an increased effort to outsource nonessential elements from one company network to that of a service provider network. Moreover, there is a movement to coordinate resource sharing between resources that are not subject to the same management system, but still address issues of security, policy, payment, and membership. For example, resources on an individual's desktop are not typically subject to the same management system as resources of a company server cluster. Even different administrative groups within a company network may implement distinct management systems.

[0005] The problems with decentralizing the resources available from servers and other computing systems operating on different network platforms, located in different regions, with different security protocols and each controlled by a different management system, have led to the development of Grid technologies using open standards for operating a grid environment. Grid environments support the sharing and coordinated use of diverse resources in dynamic, distributed, virtual organizations. A virtual organization is created within a grid environment when a selection of resources from geographically distributed systems operated by different organizations with differing policies and management systems is organized to handle a job request.

[0006] An important attribute of a grid environment that distinguishes a grid

environment from merely that of another management system is quality of service maintained across multiple diverse sets of resources. A grid environment preferably does more than just provide resources; a grid environment provides resources with a particular level of service including response time, throughput, availability, security, and the co-allocation of multiple
5 resource types to meet complex user demands. In an effort to provide quality of service, however, the issue in a grid environment is how to meet performance requirements when the reality of network systems is that optimal performance is not always available.

[0007] First, a reality of network systems is that applications are typically written to execute on specific platforms with specific operational requirements. In particular, the
10 operational requirements of applications are often specified from measurements run under optimal conditions on the particular platform. Thus, application behavior differs extensively when run on a non-native platform or resources. Therefore, when multiple heterogeneous systems are linked together in a grid environment, there is an issue of how to maintain quality of service when applications are executing on non-native platforms. In particular, there is currently
15 no means to express each application's operational requirements for execution in a grid environment where the available resources with which to execute a job may change rapidly over an interval of time.

[0008] Further a reality of network systems, is that applications are typically written where modules of the application are written to execute on specific platforms with relatively
20 static resource pools. Thus, when an application written for a relatively static resource pool is encounters suboptimal operating conditions, the entire application shuts down.

[0009] Moreover, within a grid environment, applications are just one level of the

functionality of a grid architecture. In addition to the application level, multiple levels and types of services are available to be implemented by an application. The same issues for maintaining quality of service for applications apply when these service layers run on multiple heterogeneous platforms within a grid environment.

5 **[0010]** Therefore, in view of the foregoing, there is a need for a method, system, and program to manage the software layers of functionality within a grid environment and in particular, to change an application or service behavior to maintain quality of service. Thus, there is a need for a method, system, and program to express an application or service's operational requirements for use in a grid environment, such that when suboptimal performance
10 is detected in the grid environment, the operational behavior of an application or service can be reconfigured based on the operational requirements of the application.

SUMMARY OF THE INVENTION

[0011] In view of the foregoing, the present invention relates is general to grid environments and in particular provides a method, system, and program for managing applications in a grid environment. Still more particularly, the present invention relates to maintaining application performance within a suboptimal grid environment by reconfiguring the application according to an application profile which expresses the operational requirements of an application in a grid environment.

[0012] In one embodiment, an application is submitted to at least one resource node from among multiple resource nodes within a grid environment. Then, a service availability management agent monitors a performance status of the at least one resource node. The service availability management agent compares the performance status with an operational requirement specified for when the application is operating at the at least one resource node. In particular, the operational requirement is specified for the platform on which the at least one resource node is positioned. Further, a profile for the application designates specifies the operational requirements for the application across multiple types of platforms. If the performance status does not meet the operational requirement, then the service availability management agent adjust the use by the application of the at least one resource node and other resource nodes of the grid environment, such that the application continues to operate when suboptimal conditions arise in a grid environment.

[0013] The service availability management agent may adjust both the type and amount of use of resources by the application. First, the service availability management agent may locate an alternate resource node that meets the operational requirements specified for the application for the type of platform at which the resource node is positioned and relocate the application to the alternate resource node. Second, the service availability management agent may identify a first module of the application from the application profile and send an instruction for the at least one resource node to shutdown the first module. In particular, modules are independent processing units coordinated through the application. After shutting down a module, the service availability management agent determines whether the application can still function and, if the application can still function monitors the resource nodes to determine whether additional modules need to be shutdown.

[0014] In general, the goal of the application profile is to express the operational requirements of the application across multiple heterogeneous resource platforms and to express the priority of modular breakdown of an application so that usage of resources by an application can be adjusted when suboptimal conditions are detected for the application. In one embodiment, the application profile is an XML schema that designates the attributes of the application and the performance range for the application when operating on different platforms and when shutting down modules of the application. In an alternate embodiment, the application profile may be implemented through other protocol schemas and data files.

[0015] All embodiments, features, and advantages of the present invention will become

apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further
5 objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

[0017] **Figure 1** depicts one embodiment of a computer system which may be
10 implemented in a grid environment and in which the present invention may be implemented;

[0018] **Figure 2** is block diagram illustrating one embodiment of the general types of components within a grid environment;

15 [0019] **Figure 3** is a block diagram depicting one example of an architecture that may be implemented in a grid environment;

[0020] **Figure 4** is a block diagram depicting a system operating within the grid environment from the perspective of the grid management system is depicted in accordance with
20 the method, system, and program of the present invention;

[0021] **Figure 5** is an illustrative representation depicting one embodiment of the

logical infrastructure of a grid environment in which the present invention may be implemented;

[0022] Figure 6 is a block diagram depicting one embodiment of a SAMA controller in accordance with the method, system, and program of the present invention;

5

[0023] Figure 7 is a block diagram illustrating one method of adjusting use of resources by rerouting an application in accordance with the method, system, and program of the present invention

10

[0024] Figure 8 depicts a block diagram illustrating another method of adjusting use of resources by reconfiguring an application's operational behavior in accordance with the method, system, and program of the present invention;

15

[0025] Figures 9A-9B depict a high level logic flowchart of a process and program for adjusting the use of grid resources by an application within a grid environment operating at suboptimal conditions; and

[0026] Figure 10 depicts a high level logic flowchart of a process and program for handling status requests in a grid environment.

20

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0027] Referring now to the drawings and in particular to **Figure 1**, there is depicted one embodiment of a computer system which may be implemented in a grid environment and in which the present invention may be implemented. As will be further described, the grid environment includes multiple computer systems managed to provide resources. Additionally, as will be further described, the present invention may be executed in a variety of computer systems, including a variety of computing systems, mobile systems, and electronic devices operating under a number of different operating systems managed within a grid environment.

[0028] In one embodiment, computer system **100** includes a bus **122** or other device for communicating information within computer system **100**, and at least one processing device such as processor **112**, coupled to bus **122** for processing information. Bus **122** preferably includes low-latency and higher latency paths that are connected by bridges and adapters and controlled within computer system **100** by multiple bus controllers. When implemented as a server system, computer system **100** typically includes multiple processors designed to improve network servicing power.

[0029] Processor **112** may be a general-purpose processor such as IBM's PowerPC™ processor that, during normal operation, processes data under the control of operating system and application software accessible from a dynamic storage device such as random access memory (RAM) **114** and a static storage device such as Read Only Memory (ROM) **116**. The operating system may provide a graphical user interface (GUI) to the user. In a preferred embodiment, application software contains machine executable instructions that when executed

on processor **112** carry out the operations depicted in the flowcharts of **Figures 9A-9B, 10**, and others described herein. Alternatively, the steps of the present invention might be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

5 **[0030]** The present invention may be provided as a computer program product, included on a machine-readable medium having stored thereon the machine executable instructions used to program computer system **100** to perform a process according to the present invention. The term “machine-readable medium” as used herein includes any medium that participates in providing instructions to processor **112** or other components of computer system

10 **100** for execution. Such a medium may take many forms including, but not limited to, non-volatile media, volatile media, and transmission media. Common forms of non-volatile media include, for example, a floppy disk, a flexible disk, a hard disk, magnetic tape or any other magnetic medium, a compact disc ROM (CD-ROM) or any other optical medium, punch cards or any other physical medium with patterns of holes, a programmable ROM (PROM), an

15 erasable PROM (EPROM), electrically EPROM (EEPROM), a flash memory, any other memory chip or cartridge, or any other medium from which computer system **100** can read and which is suitable for storing instructions. In the present embodiment, an example of a non-volatile medium is mass storage device **118** which as depicted is an internal component of computer system **100**, but will be understood to also be provided by an external device. Volatile media

20 include dynamic memory such as RAM **114**. Transmission media include coaxial cables, copper wire or fiber optics, including the wires that comprise bus **122**. Transmission media can also take the form of acoustic or light waves, such as those generated during radio frequency or

infrared data communications.

[0031] Moreover, the present invention may be downloaded as a computer program product, wherein the program instructions may be transferred from a remote virtual resource, such as a virtual resource **160**, to requesting computer system **100** by way of data signals

5 embodied in a carrier wave or other propagation medium via a network link **134** (e.g. a modem or network connection) to a communications interface **132** coupled to bus **122**. Virtual resource **160** may include a virtual representation of the resources accessible from a single system or systems, wherein multiple systems may each be considered discrete sets of resources operating on independent platforms, but coordinated as a virtual resource by a grid manager.

10 Communications interface **132** provides a two-way data communications coupling to network link **134** that may be connected, for example, to a local area network (LAN), wide area network (WAN), or an Internet Service Provider (ISP) that provide access to network **102**. In particular, network link **134** may provide wired and/or wireless network communications to one or more networks, such as network **102**, through which use of virtual resources, such as virtual resource

15 **160**, is accessible as provided by a grid management system **150**. Grid management system **150** may be part of multiple types of networks, including a peer-to-peer network, or may be part of a single computer system, such as computer system **100**.

[0032] As one example, network **102** may refer to the worldwide collection of networks and gateways that use a particular protocol, such as Transmission Control Protocol (TCP) and

20 Internet Protocol (IP), to communicate with one another. Network **102** uses electrical, electromagnetic, or optical signals that carry digital data streams. The signals through the various networks and the signals on network link **134** and through communication interface **132**,

which carry the digital data to and from computer system **100**, are exemplary forms of carrier waves transporting the information. It will be understood that alternate types of networks, combinations of networks, and infrastructures of networks may be implemented.

[0033] When implemented as a server system, computer system **100** typically includes
5 multiple communication interfaces accessible via multiple peripheral component interconnect (PCI) bus bridges connected to an input/output controller. In this manner, computer system **100** allows connections to multiple network computers.

[0034] Additionally, although not depicted, multiple peripheral components and internal/external devices may be added to computer system **100**, connected to multiple
10 controllers, adapters, and expansion slots coupled to one of the multiple levels of bus **122**. For example, a display device, audio device, keyboard, or cursor control device may be added as a peripheral component.

[0035] Those of ordinary skill in the art will appreciate that the hardware depicted in **Figure 1** may vary. Furthermore, those of ordinary skill in the art will appreciate that the
15 depicted example is not meant to imply architectural limitations with respect to the present invention.

[0036] With reference now to **Figure 2**, a block diagram illustrates one embodiment of the general types of components within a grid environment. In the present example, the
20 components of a grid environment **240** include a client system **200** interfacing with a grid management system **150** which interfaces with server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, and networks **230**. For purposes of illustration, the

network locations and types of networks connecting the components within grid environment **240** are not depicted. It will be understood, however, that the components within grid environment **240** may reside atop a network infrastructure architecture that may be implemented with multiple types of networks overlapping one another. Network infrastructure may range from multiple large enterprise systems to a peer-to-peer system to a single computer system. Further, it will be understood that the components within grid environment **240** are merely representations of the types of components within a grid environment. A grid environment may simply be encompassed in a single computer system or may encompass multiple enterprises of systems.

[0037] The central goal of a grid environment, such as grid environment **240** is organization and delivery of resources from multiple discrete systems viewed as virtual resource **160**. Client system **200**, server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, networks **230** and the systems creating grid management system **150** may be heterogeneous and regionally distributed with independent management systems, but enabled to exchange information, resources, and services through a grid infrastructure enabled by grid management system **150**. Further, server clusters **222**, servers **224**, workstations and desktops **226**, data storage systems **228**, and networks **230** may be geographically distributed across countries and continents or locally accessible to one another.

[0038] In the example, client system **200** interfaces with grid management system **150**.

Client system **200** may represent any computing system sending requests to grid management system **150**. In particular, client system **200** may send job requests and jobs to grid management system **150**. Further, while in the present embodiment client system **200** is depicted as accessing

grid environment **240** with a request, in alternate embodiments client system **200** may also operate within grid environment **240**.

[0039] While the systems within virtual resource **160** are depicted in parallel, in reality, the systems may be part of a hierarchy of systems where some systems within virtual resource **160** may be local to client system **200**, while other systems require access to external networks. Additionally, it is important to note, that systems depicted within virtual resources **160** may be physically encompassed within client system **200**.

[0040] One function of grid management system **150** is to manage job requests and jobs from client system **200** and control distribution of each job to a selection of computing systems of virtual resource **160** for use of particular resources at the available computing systems within virtual resource **160**. From the perspective of client system **200**, however, virtual resource **160** handles the request and returns the result without differentiating between which computing system in virtual resource **160** actually performed the request.

[0041] To implement grid environment **240**, grid management system **150** facilitates grid services. Grid services may be designed according to multiple architectures, including, but not limited to, the Open Grid Services Architecture (OGSA). In particular, grid management system **150** refers to the management environment which creates a grid by linking computing systems into a heterogeneous network environment characterized by sharing of resources through grid services.

[0042] Grid environment **240**, as managed by grid management system **150**, may provide a single type of service or multiple types of services. For example, computational grids, scavenging grids, and data grids are example categorizations of the types of services provided in

a grid environment. Computational grids may manage computing resources of high-performance servers. Scavenging grids may scavenge for CPU resources and data storage resources across desktop computer systems. Data grids may manage data storage resources accessible, for example, to multiple organizations or enterprises. It will be understood that a grid environment
5 is not limited to a single type of grid categorization.

[0043] Referring now to **Figure 3**, a block diagram illustrates one example of an architecture that may be implemented in a grid environment. As depicted, an architecture **300** includes multiple layers of functionality. As will be further described, the present invention is a
10 process which may be implemented in one or more layers of an architecture, such as architecture **300**, which is implemented in a grid environment, such as the grid environment described in **Figure 2**. It is important to note that architecture **300** is just one example of an architecture that may be implemented in a grid environment and in which the present invention may be implemented. Further, it is important to note that multiple architectures may be implemented
15 within a grid environment.

[0044] Within architecture **300**, first, a physical and logical resources layer **330** organizes the resources of the systems in the grid. Physical resources include, but are not limited to, servers, storage media, and networks. The logical resources virtualize and aggregate the physical layer into usable resources such as operating systems, processing power, memory, I/O
20 processing, file systems, database managers, directories, memory managers, and other resources.

[0045] Next, a web services layer **320** provides an interface between grid services **310** and physical and logical resources **330**. Web services layer **320** implements service interfaces

including, but not limited to, Web Services Description Language (WSDL), Simple Object Access Protocol (SOAP), and eXtensible mark-up language (XML), executing atop an Internet Protocol(IP) or other network transport layer. Further, the Open Grid Services Infrastructure (OSGI) standard **322** builds on top of current web services **320** by extending web services **320** to
5 provide capabilities for dynamic and manageable Web services required to model the resources of the grid. In particular, by implementing OGSi standard **322** with web services **320**, grid services **310** designed using OGSA are interoperable. In alternate embodiments, other infrastructures or additional infrastructures may be implemented a top web services layer **320**.

[0046] Grid services layer **310** includes multiple services. For example, grid services
10 layer **310** may include grid services designed using OGSA, such that a uniform standard is implemented in creating grid services. Alternatively, grid services may be designed under multiple architectures. Grid services can be grouped into four main functions. It will be understood, however, that other functions may be performed by grid services.

[0047] First, a resource management service **302** manages the use of the physical and
15 logical resources. Resources may include, but are not limited to, processing resources, memory resources, and storage resources. Management of these resources includes receiving job requests, scheduling job requests, distributing jobs, and managing the retrieval of the results for jobs. Resource management service **302** preferably monitors resource loads and distributes jobs to less busy parts of the grid to balance resource loads and absorb unexpected peaks of activity.
20 In particular, a user may specify preferred performance levels so that resource management service **302** distributes jobs to maintain the preferred performance levels within the grid.

[0048] Second, information services **304** manages the information transfer and

communication between computing systems within the grid. Since multiple communication protocols may be implemented, information services **304** preferably manages communications across multiple networks utilizing multiple types of communication protocols.

[0049] Third, a data management service **306** manages data transfer and storage within
5 the grid. In particular, data management service **306** may move data to nodes within the grid where a job requiring the data will execute. A particular type of transfer protocol, such as Grid File Transfer Protocol (GridFTP), may be implemented.

[0050] Finally, a security service **308** applies a security protocol for security at the connection layers of each of the systems operating within the grid. Security service **308** may
10 implement security protocols, such as Open Secure Socket Layers (SSL), to provide secure transmissions. Further, security service **308** may provide a single sign-on mechanism, so that once a user is authenticated, a proxy certificate is created and used when performing actions within the grid for the user.

[0051] Multiple services may work together to provide several key functions of a grid
15 computing system. In a first example, computational tasks are distributed within a grid. Data management service **306** may divide up a computation task into separate grid services requests of packets of data that are then distributed by and managed by resource management service **302**. The results are collected and consolidated by data management system **306**. In a second example, the storage resources across multiple computing systems in the grid are viewed as a
20 single virtual data storage system managed by data management service **306** and monitored by resource management service **302**.

[0052] An applications layer **340** includes applications that use one or more of the grid

services available in grid services layer **310**. Advantageously, applications interface with the physical and logical resources **330** via grid services layer **310** and web services **320**, such that multiple heterogeneous systems can interact and interoperate. As an example, an application may be available at a web site “www.ibm.com”. The web site provides an entry point for the user to send a job request to select from services, such as a catalog search engine or a business to business service. It is important to note that while an application and a grid service are positioned within different layers of the grid architecture **300**, a reference to an application in general also encompasses grid services exposed through the application.

[0053] With reference now to **Figure 4**, a block diagram of a system operating within the grid environment from the perspective of the grid management system is depicted in accordance with the method, system, and program of the present invention. As illustrated in **Figure 1**, a computer system includes many hardware components. As part of a grid environment, however, these hardware components are viewed as resources. For example, a system **400** includes an application resource **402**, two CPU resources **404** and **406**, a memory resource **408**, and a storage resource **410**. The resources in system **400** are typical of the types of resources when viewed within the grid environment, however, in an alternate embodiment, other types of resources may be provided. Further, the resources in system **400** may be physically located within a single computer system or distributed across multiple computer systems connected by a network, for example.

[0054] As part of the grid management system described in **Figure 2**, a grid manager and router (GM) **424** provides the interface between the resources of system **400** and client

systems sending requests.

[0055] In particular, a resource monitor **422** within GM **424** monitors the working status of each of the resources available in system **400**. GM **424** preferably sends status reports to other grid managers and routers within the grid environment to indicate the availability of the resources in system **400**. Additionally, status reports may describe the computer hardware, operating system, and resources of system **400**. Status reports may be generated, for example, when system **400** joins or leaves the grid environment, when a threshold is detected, at predetermined time intervals, and on specific predetermined events, including, but not limited to a hardware fault or a portion of an application or service failing.

[0056] Referring now to **Figure 5**, an illustrative representation depicts one embodiment of the logical infrastructure of a grid environment in which the present invention may be implemented. While **Figure 2** depicts an example of general components of a grid environment, in the present figure, an example of how the general components are viewed logically within a grid environment is illustrated in a grid environment **540**. In particular, the grid management system functions are logically dispersed into multiple GMs, such as GM **504**, GM **510**, and GM **520**. Further, the virtual resource is logically dispersed into multiple resources (RSs), such as RS **506**, RS **508**, RS **512**, RS **514**, RS **522**, and RS **524**. It is important to note that a resource may not be a direct representation of a physical resource, but rather a logical representation of one or more physical resources and or groups of physical resources.

[0057] In the example, client system **200** sends a job request to GM **504**. GM **504** searches for resources available to handle the job specified in the job request. In particular, GM

504 checks whether **RS 506** and **RS 508** can handle the job specified in the job request and may send queries to other GMs, such as **GM 510** or **GM 520**. GMs **510** and **520** return reports on the availability of resources to handle the job request.

[0058] In particular, a job request preferably includes a request made through a particular application for a grid service. In general, an application is self-contained, but exposes the grid services layer. The resources necessary to handle the job request are those necessary to handle the request as specified by the application and services accessed.

[0059] For purposes of illustrations, **RS 506** and **RS 508** are considered local resources or resources within the same discrete set of resources to which jobs from client system **200** are submitted. In the examples following, when **RS 506** and **508** are not meeting performance requirements for a job from client system **200**, then additional resources may be allocated including other resources within the same discrete set of resources, capacity on demand resources, resources from internal grids and finally resources from external grids.

[0060] More specifically, in the example, **GM 510**, **RS 512**, and **RS 514** are part of one grid infrastructure “grid A” operated by a first business that provides a first specified number of grid services for a first specified price. Then, **GM 520**, **RS 522**, and **RS 524** are part of another grid infrastructure “grid B” operated by a second business that provides a second specified number of grid services for a second specified price. When **GM 504** sends the job request to **GM 510** and **GM 520**, the each GM preferably reports whether the job request can be handled and a price for handling the request. In relation to client system **200**, grids A and B may be internal grids operating within the same enterprise system as client system **200** or external grids.

[0061] After receiving reports on the availability of resources, **GM 504** collects the

options for handling the job and returns the options to client system **200**. Client system **200** may then decide to select a preferred option and send the job to have handled according to the selected option. GM **504** manages the return of the results of the job to client system **200**.

[0062] The resources utilized in the option selected by client system **200** form a virtual
5 organization for handling the job. For example, if client system **200** selects an option to handle the job based on the option returned from GM **510**, then a virtual organization may be formed to process the job which would include GM **504**, GM **510**, RS **512**, and RS **514**.

[0063] According to an advantage of the present invention, a service availability
management agent (SAMA) **530** is preferably accessible within grid environment **540**. SAMA
10 performs the functions of monitoring grid resources, policy coordination, application profile management, analytical processing, and problem dispatch for grid environment **540**. SAMA **530** may be controlled by dynamic policies that manage its behavior while performing error management. In particular, errors may occur within grid environment **530** when suboptimal conditions occur that degrade applications and services to a minimal operational level or to the
15 point of no service at all. Such suboptimal performance may occur as a result of a system failure, a network infrastructures dropping or becoming overloaded, or other failures that occur within grid environment **540**.

[0064] In particular, when suboptimal conditions are detected within grid environment **540**, SAMA **530** preferably manages use of grid resources to allow applications and services to
20 continue to function. In particular, SAMA **530** may regulate the type and amount of resources provided to an application or service in an attempt keep the application or service running. For example, SAMA **530** may move an application or service to a new set of resource nodes.

Additionally, for example, SAMA **530** may reconfigure an application's operation behavior to use less resources.

[0065] In particular, SAMA **530** may be incorporated within a single system or distributed across multiple systems. Further, within grid architecture **300** of **Figure 3**, SAMA **530** may be implemented in multiple levels. For example, SAMA **530** may be implemented in part as a web service and in part as a resource management service.

[0066] It is important to note that while SAMA **530** may manage use of grid resources by both applications and services, for purposes of example, management of grid resources for applications will be described. Further, it will be understood that SAMA **530** may manage use of grid resources by other controllers and agents operating within grid environment **540**.

[0067] With reference now to **Figure 6**, there is a block diagram of one embodiment of a SAMA controller in accordance with the method, system, and program of the present invention. As depicted, SAMA **530** includes multiple controllers and a SAMA database **610**.

[0068] First, a resource monitor **606** monitors the performance and availability of resources in a grid environment. More specifically, resource monitor **606** may monitor resource usage, capacity, and throughput, for example.

[0069] An inference analysis controller **604** preferably receives messages from systems operating in the grid environment and parses the messages to determine the specific request. In particular, a message may include a job request. A job request may invoke a particular service or services published for an application available from the grid. Each application preferably has a profile stored in profiles **612**. For purposes of example, grid application profiles database **612**

includes Document Type Definitions (DTDs) of an XML expression of the profile of each application. It is important to note that while the examples of the application of profiles to manage resource usage are focused on resource usage by applications, a profile may be specified for all layers of functionality within the grid architecture, including services.

- 5 **[0070]** An example schema of an application profile is illustrated in **Table 1**. In the example, a sample of the types of application attribute information that may be included in a application profile is illustrated. In particular, the name, version, description, developer name, owner name, and code size for an application may be specified. Next, in the example, a sample of the minimal performance and maximum performance requirements of an application are
- 10 specified for each platform available for operating the application within the grid environment. Further, modules of the application and expected performance based on platforms are included. A module definition may include, for example, the module name, module data, amount of resources required and priority number. A platform definition may include, for example, a name of the platform, version and description of the platform, maximum and minimum memory size,
- 15 maximum and minimum CPU requirements on a symmetrical multi-processor (SMP) system, maximum and minimum CPU clock speeds, and the operating name and version. In addition, as a security requirement, for any platform on which the application will operate, the application profile may specify a particular type of required platform certificate where the certificate indicates that the platform has under gone some level of assurances by a security standard, such
- 20 as Common Criteria (CC). It will be understood that while the present example of the application profile is an XML expression, other types of profile expressions may be implemented. It will be understood that the example schema may also applied to services, agents, and other controllers

within a grid environment. Further, as will be understood by those with skill in the art, other types of information may be included in an application profile in an implementation of the present invention.

[0071] Profile management controller **602** compares the maximum and minimum
5 operating requirements of an application profile to the current operating conditions to determine if the grid is operating at optimal conditions. If the grid is not running at optimal conditions, then profile management controller **602** preferably determines whether the application can be scheduled to run at another resource node or nodes. In particular, profile management controller **602** compares the platform of the additional resource nodes with the platform requirements of
10 the application as specified in the application profile. Further, if additional nodes are not available, then profile management controller **602** determines whether modules of the application can be degraded or shutdown, while still maintaining minimum operating conditions according to the application profile. In particular, profile management controller **602** may process the application profile to determine the modules identified for the application. In one
15 example, each module is assigned a resource size requirement and priority to be pruned or shut down. Profile management controller **602** preferably shuts down modules according to each module's priority and resource size requirements.

[0072] Alternatively, a message received at SAMA **530** may include a status request for the availability of the grid environment to handle a job request before a client system sends a job.
20 In particular, a service inventory database **616** maintains the current status of applications and services available from the grid. Service inventory database **616** is preferably updated whenever modules are shutdown, such that responses to status requests are efficiently returned based on

the status in service inventory database **616**.

[0073] Before profile management controller **602** adjusts the use of grid resources, inference analysis controller **604** may compare the application request or monitored suboptimal condition with grid policies in a policy database **614**. In particular, a grid policies may specify overall operational requirements for the grid environment. Further, grid policies may specify the type and amount of resources that are available for a particular application when profile management controller **602** is adjusting the use of grid resources.

Table 1

<?xml version="1.0" encoding="UTF-8"?>

<!--

Application Profile DTD - Version 1.0

+:One or more permitted

*:Zero or more permitted

?:Optional

-->

<!-- Application Profile Definition -->

<!ELEMENT Application(ApplicationData, AppVital Security+)>

<!--Application attribute information -->

5

<!ELEMENT ApplicationData >

<!--ATTLIST ApplicationData

Name CDATA # REQUIRED

Version CDATA # REQUIRED

10 Description CDATA #REQUIRED

DeveloperName CDATA #REQUIRED

OwnerName CDATA #REQUIRED

codeSize CDATA #REQUIRED

>

15

<!ELEMENT AppVital (PruneModule*, Platform+)>

<!--ATTLIST AppVital

minimalPerformance CDATA #REQUIRED

maximumPerformance CDATA #REQUIRED

20

<!-- Prune Module Definition -->

<![ELEMENT PruneModule EMPTY]>

<![ATTLIST PruneModule

ModuleName CDATA #REQUIRED

5 ModuleData CDATA #REQUIRED

ResourceAmt CDATA #REQUIRED

PriorityNumber CDATA #REQUIRED >

<!-- Platform supported Definition -->

10

<![ELEMENT Platform EMPTY]>

<![ATTLIST Platform

Name CDATA #REQUIRED

Version CDATA #REQUIRED

15 Description CDATA #REQUIRED

MaxMemorySize CDATA #REQUIRED

MinMemorySize CDATA #REQUIRED

MaxCPU CDATA #REQUIRED

MinCPU CDATA #REQUIRED

20 MaxSpeed CDATA #REQUIRED

Min CDATA #REQUIRED

OSName CDATA #REQUIRED

OSVersion CDATA #REQUIRED

evalName CDATA #IMPLIED

evalProvider CDATA #IMPLIED

evalLevel CDATA #IMPLIED

5 evalDateIssued CDATA #IMPLIED

>

<!-- Security Information Definition -->

10 <!--ELEMENT Security EMPTY-->

<!--ATTLIST Security

keyType (SYMMETRIC|ASYMMETRIC) CDATA # REQUIRED

algorithmName CDATA # REQUIRED

algorithmAttribute CDATA #IMPLIED

15 keyLength CDATA #REQUIRED

keyIdentifier CDATA #REQUIRED

keyUsage (Authentication|Encryption|Signature) CDATA #REQUIRED

eExportable (TRUE|FALSE) CDATA #REQUIRED

certificateType (X509V1|X509V3|OTHER) CDATA #IMPLIED

20 certificateName CDATA #IMPLIED

certificateAttribute CDATA #IMPLIED

cSPProviderType (Software|Hardware) CDATA #REQUIRED

ssmName CDATA #IMPLIED

ssmVersion CDATA #IMPLIED

ssmExtName CDATA #IMPLIED

ssmExtVersion CDATA #IMPLIED

5 hsmName CDATA #IMPLIED

hsmVersion CDATA #IMPLIED

hsmFirmName CDATA #IMPLIED

hsmFirmVersion CDATA #IMPLIED

hsmExtName CDATA #IMPLIED

10 hsmExtVersion CDATA #IMPLIED

evalName CDATA #IMPLIED

evalProvider CDATA #IMPLIED

evalLevel CDATA #IMPLIED

evalDateIssued CDATA #IMPLIED

15 >

[0074] With reference now to **Figure 7**, there is depicted a block diagram illustrating one method of adjusting use of resources by rerouting an application in accordance with the

20 method, system, and program of the present invention. As depicted, an application **700** is initially run on a resource at node A operating on platform **706**. The operational requirements for application **700** to run on platform **706** are preferably defined in a profile for application **700**.

If the performance of application **700** running on platform **706** is suboptimal according to the profile, then the SAMA may search for other resources available to meet the operational requirements of application **700**. In the example, a node B running on platform **708** is available to meet the operational requirements specified for platform **708** by the application profile for application **700**. Once a new node is located for application **700**, the SAMA may direct a job router to redirect the application from resource node A to node B and begin monitoring the performance of node B. Although the present example is described with reference to an application moving from one node on one platform to another node on another platform, it will be understood that an application may move from one resource node to another within the same platform or may move from one cluster of resource nodes to another cluster of resource nodes.

[0075] Referring now to **Figure 8**, there is depicted a block diagram illustrating another method of adjusting use of resources by reconfiguring an application's operational behavior in accordance with the method, system, and program of the present invention. As illustrated, application **700** includes multiple independent processing modules A-N. Modules A-N may be handled by a single resource or multiple resources. Further, a single module may be handled across multiple resources, depending on the size of resources required for the module.

[0076] In particular, in the profile for application **700**, a shutdown priority is preferably specified for each module. In the example, modules A, B and N of application **700** are executing in the grid environment on one or more nodes. The grid environment is only able to provide suboptimal performance for application **700**. Thus, the SAMA determines whether modules of application **700** can be shutdown so that portions of application **700** can continue to operate

without the entire application shutting down. In the example, module B is assigned the highest shutdown priority. The SAMA may determine that this highest shutdown priority module consumes resources to the extent that the module should be shutdown. The SAMA may send an instruction to the node or nodes running module B to shutdown the module. If the performance
5 remains suboptimal after shutting down module B, the SAMA may send instructions to shutdown additional modules.

[0077] With reference now to **Figures 9A-9B**, there is illustrated a high level logic flowchart of a process and program for adjusting the use of grid resources by an application
10 within a grid environment operating at suboptimal conditions. As illustrated, the process starts at block **900** and thereafter proceeds to block **902**. Block **902** depicts a determination whether a job request is received. If a job request is not received, then the process iterates at block **902**. If a job request is received, then the process passes to block **904**. Block **904** depicts retrieving the application profile of at least one application invoked through the application from which the job
15 request originated. Next, block **905** depicts scheduling resources according to the application profile and grid policies. Thereafter, block **906** depicts submitting the job to the selected grid resources, and the process passes to block **907**.

[0078] Block **907** illustrates monitoring grid performance of the resources running an application. Thereafter, block **908** depicts a determination whether suboptimal conditions are
20 detected for the application when compared with the application profile. If suboptimal conditions are not detected, then the process returns to block **907**. If suboptimal conditions are detected, then the process passes to block **910**.

[0079] Block 910 depicts searching for other nodes to relocate the application to meet the application operational requirements from the application profile for the particular platform(s) on which the application is executing. Next, block 912 depicts a determination whether relocation nodes are available. This determination may first require determining, based on grid policies, which resources are available for the application. Then, a determination may be made of which of the available resources meet the operational requirements of the application profile as specified for platforms supporting the available resources. If relocation nodes are available, then block 920 depicts relocating the application to the relocation nodes, and the process returns to block 907 while the application is invoked. If relocation nodes are not available, then the process passes to block 914. Block 914 depicts analyzing the application profile to determine a first priority module of the application. Next, block 916 depicts a determination whether SAMA should shut down the module. In particular, SAMA may determine if shutting down the module would free sufficient resources to increase overall performance of the application. Additionally, grid policies and additional application profile information may control whether the SAMA decides to shut down the module. If SAMA should not shut down the module, then the process passes to block 932, to be further described. If SAMA should shut down the module, then the process passes to block 918.

[0080] Block 918 depicts sending an XML message to the node running the application where the message includes authentication and authorization to reconfigure the application. It will be understood that other transaction protocols may be implemented. Next, block 922 depicts a determination whether an acknowledgment that the module was shut down is received. If the acknowledgement is not received, the process iterates at block 922 for a period of time. If

the acknowledgement is received, then the process passes to block **924**. Block **924** depicts reevaluating the grid status based on the resource performance. Next, block **926** depicts a determination whether additional steps are required to reconfigure the application. In particular, the application may need additional reconfiguration if after shutting down one module the application continues to operate under suboptimal conditions. If additional steps are required, then the process passes to block **932**, to be further described. If additional steps are not required, then the process passes to block **928**. Block **928** depicts allowing the application to run with degraded modules. Next, block **930** depicts updating the inventory of available applications with the degraded level, and the process ends.

10 **[0081]** As previously references, the process may passes to block **932**. Block **932** depicts a determination whether the application can still function under a degraded state. If the application cannot still function, then the process passes to block **936**. Block **936** depicts alerting the requester with an error message, and the process ends. Alternatively, if the application can still function, then the process passes to block **934**. Block **934** depicts analyzing
15 the application profile to determine a next priority module of the application, and the process passes to block **916**.

[0082] Referring now to **Figure 10**, there is depicted a high level logic flowchart of a process and program for handling status requests in a grid environment. As illustrated, the
20 process starts at block **1000** and thereafter proceeds to block **1002**. Block **1002** depicts receiving a request to pre-determine the sufficiency of the resources for an application. Next, block **1004** depicts checking the service inventory database for current sufficiency of resources. A monitor

request may be sent to the potentially effected resources to determine the current availability of resources. Finally, block 1006 depicts returning an XML status request response of the sufficiency of resources, and the process ends.

5 **[0083]** While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention. In particular, it is important to note that while the description of the present invention focuses on the invention operating in the context of grid applications, services, agents, and
10 controllers, the present invention may also apply in the context of normal web services and in distributed architectures in general.